arXiv:1406.4549v1 [stat.ME] 17 Jun 2014

Extensible grids: uniform sampling on a space-filling curve

Zhijian He
Tsinghua University

Art B. Owen
Stanford University

June 2014

Abstract

We study the properties of points in $[0,1]^d$ generated by applying Hilbert's space-filling curve to uniformly distributed points in $[0,1]$. For deterministic sampling we obtain a discrepancy of $O(n^{-1/d})$ for $d \geq 2$. For random stratified sampling, and scrambled van der Corput points, we get a mean squared error of $O(n^{-1-2/d})$ for integration of Lipshitz continuous integrands, when $d \geq 3$. These rates are the same as one gets by sampling on d dimensional grids and they show a deterioration with increasing d. The rate for Lipshitz functions is however best possible at that level of smoothness and is better than plain IID sampling. Unlike grids, space-filling curve sampling provides points at any desired sample size, and the van der Corput version is extensible in n. Additionally we show that certain discontinuous functions with infinite variation in the sense of Hardy and Krause can be integrated with a mean squared error of $O(n^{-1-1/d})$. It was previously known only that the rate was $o(n^{-1})$. Other space-filling curves, such as those due to Sierpinski and Peano, also attain these rates, while upper bounds for the Lebesgue curve are somewhat worse, as if the dimension were $\log_2(3)$ times as high.

Keywords: Hilbert space-filling curve, Lattice sequence, van der Corput sequence, randomized quasi-Monte Carlo. sequential quasi-Monte Carlo.

1 Introduction

A Hilbert curve is a continuous mapping $H(x)$ from $[0,1]$ to $[0,1]^d$ for $d > 1$. It is an example of a class of space-filling curves, of which Peano's was first. Space-filling curves have long been mathematically intriguing, but since the 1980s (see Bader (2013)) they have become important computational tools in computer graphics, in finding near optimal solutions to the travelling salesman problem, and in PDE solvers where elements in a multidimensional mesh must be allocated to a smaller number of processors (Zumbusch, 2003). In this paper we look at a quasi-Monte Carlo method that takes equidistributed points $x_i \in [0,1]$ and then uses $P_i = H(x_i) \in [0,1]^d$. The analysis also provides convergence rates for some functions that are not smooth enough to benefit from unrandomized quasi-Monte Carlo sampling.

Our interest in this problem was sparked by Gerber and Chopin (2014), who present an innovative combination of sequential Monte Carlo (SMC), quasi-Monte Carlo (QMC), and Markov chain Monte Carlo (MCMC) as a method to compete with particle MCMC. The resulting method is closely related to the array-RQMC algorithm of L'Ecuyer et al. (2008).

The particle algorithms simulate N copies of a Markov chain through a sequence of time steps $t = 1, \dots, T$. At the end of time step t, chain n is in position $\boldsymbol{x}_{nt} \in \mathbb{R}^d$, $n = 1, \dots, N$. The computation to advance a chain from $\boldsymbol{x}_{nt}$ to $\boldsymbol{x}_{n,t+1}$ requires a point in $[0,1]^S$, which may either be uniformly distributed, in Monte Carlo, or from a low discrepancy ensemble, in quasi-Monte Carlo. It is possible to advance all N chains by one time step using a matrix $U^{(t+1)} \in [0,1]^{N\times(S+1)}$. The first column of $U^{(t+1)}$ is used to identify which row of $U^{(t+1)}$ will be used to advance each of the points $\boldsymbol{x}_{nt}$, and then the remaining S columns are used to advance the N chains. When $d = 1$ we can sort x_{nt} into the same order as the first column of $U^{(t+1)}$. Then if $U^{(t+1)}$ is a low discrepancy point set, the starting positions x_{nt} are equidistributed with respect to the updating variables.

Things become much more difficult when $\boldsymbol{x}_{nt} \in \mathbb{R}^d$ for $d \geq 2$. Then it is not straightforward how one should align $\boldsymbol{x}_{nt}$ with the first column or first several columns of $U^{(t+1)}$. Gerber and Chopin (2014) place a space-filling curve in $\mathbb{R}^d$. Each point $\boldsymbol{x}_{nt} \in \mathbb{R}^d$ has a coordinate on this curve, its pre-image in $[0,1]$. Then $\boldsymbol{x}_{1t}$ to $\boldsymbol{x}_{Nt}$ are sorted in increasing order of those pre-images, and the k'th largest one is aligned with the row of $U^{(t+1)}$ having the k'th largest value in column 1.

They give conditions under which their algorithm estimates expectations

with a root mean squared error of $o(n^{-1})$. Their sequential Monte Carlo scheme has a provably better rate of convergence than Monte Carlo or Markov chain Monte Carlo. Array-RQMC behaves empirically as if it has a better rate (L'Ecuyer et al., 2008) but as yet there is no proof. In principal one could simulate the chains through T steps without any remapping by using a quasi-Monte Carlo scheme in $[0,1]^{ST}$. But in such high dimensions it becomes difficult to construct point sets with meaningfully better equidistribution than Latin hypercube samples (McKay et al., 1979) have.

In this paper we examine the simpler related problem taking either a QMC or randomized QMC sample within $[0,1]$, and applying the Hilbert curve to that sample in order to get a quadrature rule in $[0,1]^d$. We study the accuracy of quadrature. This strategy has been used in computer graphics for related purposes. Rafajłowicz and Skubalska-Rafajłowicz (2008) applied a two dimensional space-filling curve to a Kronecker sequence in $[0,1]$ in order to downsample an image. They report that this strategy allows them to approximate the Fourier spectra of the images. Schretter and Niederreiter (2013) report that they can downsample images with fewer visual artifacts this way than by using a two dimensional QMC sequence.

Section 2 introduces the Hilbert curve, giving its important properties. Section 3 studies the star-discrepancy of the resulting points obtained as the d dimensional image of one-dimensional low discrepancy points. We find that the star-discrepancy is $O(n^{-1/d})$, which is very high considering that quasi-Monte Carlo rules typically attain the $O(n^{-1+\epsilon})$ rate, where $\epsilon > 0$ hides logarithmic factors. Section 4 considers some randomized quasi-Monte Carlo (RQMC) versions of Hilbert sampling. The mean squared error converges as $O(n^{-1-2/d})$ for Lipshitz continuous integrands, and as $O(n^{-1-1/d})$ for certain discontinuous integrands of infinite variation, studied there. Thus we see a better than Monte Carlo convergence rate, though one that deteriorates with increasing dimension. As a result we expect the much more complicated proposal of Gerber and Chopin (2014) to have diminishing effectiveness with increasing dimension. Section 5 presents numerical results showing a close match between mean squared error rates in our theorems and observed errors in some example functions. That is, the asymptote appears to be relevant at small sample sizes. Section 6 compares the results here to those of other methods one might use. We find that Hilbert curve quadrature commonly gives the same convergence rates that one would see from using grids of $n = m^d$ points in $[0,1]^d$, but makes those rates available at all integer sizes n. At low smoothness levels (Lipshitz continuity only) that poor rate is in fact

best possible.

2 Hilbert Curves

Here we introduce Hilbert's space-filling curve and some of its properties that we need. For more background, there is the monograph Sagan (1994) on space-filling curves, of which Chapter 2 describes Hilbert's curve. Zumbusch (2003) describes multilevel numerical methods, including Chapter 4 on space-filling curves.

Throughout this paper, d is a positive integer, λ_d is d-dimensional Lebesgue measure, and $\|\cdot\|$ is the usual Euclidean norm. For integer $m \geq 0$, define 2^{dm} intervals

$$I_d^m(k) = \left[\frac{k}{2^{dm}}, \frac{k+1}{2^{dm}}\right], \quad k = 0, \ldots, 2^{dm} - 1,$$

and let $\mathcal{I}_d^m = \left\{I_d^m(k) \mid k < 2^{dm}\right\}$. Next, for $\kappa = (k_1, \ldots, k_d)$ with $k_j \in \{0, 1, \ldots, 2^m - 1\}$ define 2^{dm} subcubes of $[0,1]^d$ via

$$E_d^m(\kappa) = \prod_{j=1}^{d} \left[\frac{k_j}{2^m}, \frac{k_j + 1}{2^m}\right]. \tag{2.1}$$

The set of indices κ is $\mathcal{K}_d^m = \{0, 1, \ldots, 2^m - 1\}^d$ and we let $\mathcal{E}_d^m = \{E_d^m(\kappa) \mid \kappa \in \mathcal{K}_d^m\}$. One can find a sequence of mappings $H_m : \mathcal{I}_d^m \to \mathcal{E}_d^m$ with the following properties,

- Bijection: For $k \neq k'$, $H_m(I_d^m(k)) \neq H_m(I_d^m(k'))$.
- Adjacency: The two subcubes $H_m(I_d^m(k))$ and $H_m(I_d^m(k+1))$ are adjacent. That is, they have one $(d-1)$-dimensional face in common.
- Nesting: If we split $I_d^m(k)$ into the 2^d successive subintervals $I_d^{m+1}(k_\ell)$, $k_\ell = 2^d k + \ell$, $\ell = 0, \ldots, 2^d - 1$, then the $H_{m+1}(I_d^{m+1}(k_\ell))$ are subcubes whose union is $H_m(I_d^m(k))$.

Figure 1 illustrates the Hilbert curve construction in dimension 2.

The Hilbert curve is defined by $H(x) = \lim_{m\to\infty} H_m(x)$. The point $x \in [0,1]$ belongs to an infinite sequence $I_d^m(k_m)$ of intervals which shrink to x. If x does not have a terminating base 2 representation then the sequence $I_d^m(k_m)$ is unique and then $H_m(I_d^m(k_m))$ is a unique sequence of subcubes.

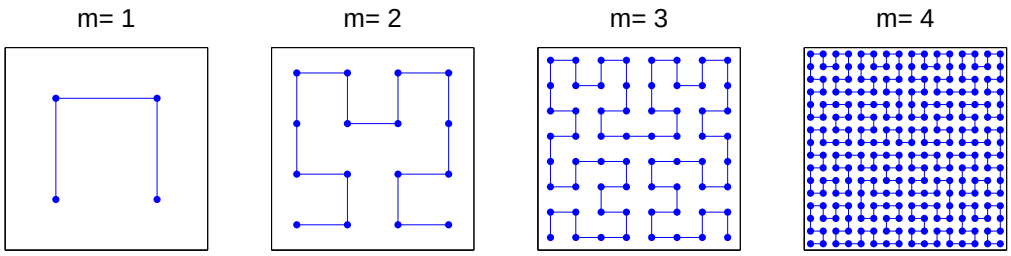

Figure 1: First 4 stages in the approximation of Hilbert's space-filling curve

Points such as $x = 1/4 = 0.01\overline{0} = 0.00\overline{1}$ with two binary representations nevertheless have uniquely defined $H(x)$. The Hilbert curve passes through every point in $[0,1]^d$. It is not surjective: there are points $x \neq x'$ with $H(x) = H(x')$. Indeed, a result of Netto (1879) shows that no space-filling curve from $[0,1]$ to $[0,1]^d$ for $d > 1$ can be bijective.

There is more than one way to define the sequence of mappings in a Hilbert curve. But any of those ways produces a mapping H with these properties:

- P(1): $H(I_d^m(k)) = H_m(I_d^m(k))$.
- P(2): If $A \subset [0,1]$ is measurable, then $\lambda_1(A) = \lambda_d(H(A))$.
- P(3): If $x \sim \mathbb{U}([0,1])$, then $H(x) \sim \mathbb{U}([0,1]^d)$. It admits the change of variables:
$$\mu = \int_{[0,1]^d} f(x)\,\mathrm{d}x = \int_0^1 f(H(x))\,\mathrm{d}x. \tag{2.2}$$
- P(4): The function $H(x)$ is Hölder continuous, but nowhere differentiable. More precisely, for any $x, y \in [0,1]$, we have
$$\|H(x) - H(y)\| \leq 2\sqrt{d+3}\,|x-y|^{1/d}, \tag{2.3}$$

The Hölder property P(4) is proved in Zumbusch (2003). We prove it here too, because the proof is short and we make extensive use of that result.

Theorem 2.1. *If $x, y \in [0,1]$ and H is Hilbert's space-filling curve in dimension $d \geq 1$, then $\|H(x) - H(y)\| \leq 2\sqrt{d+3}\,|x-y|^{1/d}$.*

Proof. Without loss of generality, $x < y$. Let $m = \lfloor -\log_2 |x-y| / d \rfloor$ so that $2^{-dm} \geq |x-y| > 2^{-d(m+1)}$. The interval $[x, y]$ is contained within one, or at most two, consecutive intervals $I_d^m(k')$, $I_d^m(k'+1)$ for some $k' < 2^{dm} - 1$. As a result, the image $H([x, y])$ lies within $H(I_d^m(k')) \cup H(I_d^m(k'+1))$. By P(1) and the adjacency property of H_m, the diameter of $H([x, y])$ is bounded by the diameter of two adjacent subcubes of the form (2.1), which is $2^{-m}\sqrt{d+3} \leq 2\sqrt{d+3}\,|x-y|^{1/d}$. □

In the context of numerical integration, the integral μ in (2.2) can be estimated by the following average:

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^{n} f(H(x_i)), \tag{2.4}$$

where x_i's are carefully chosen quadrature points in $[0, 1]$. The space-filling curve reduces a multidimensional integral to a one-dimensional numerical integration problem. It is important to point out that the integrand $f \circ H(x)$ is not of bounded variation even for smooth (but non-trivial) functions f. Bounded variation would have yielded convergence rates of $|\hat{\mu} - \mu| = O(1/n)$ in any dimension via the Koksma-Hlawka inequality (see Section 3).

3 Star-discrepancy

Given a sequence $x_1, \ldots, x_n$ in $[0, 1]$, we can obtain a corresponding sequence $P_1, \ldots, P_n$ in $[0, 1]^d$ by the Hilbert mapping function described above, $P_i = H(x_i)$.

We use the star-discrepancy to measure the uniformity of the resulting sequence $\mathcal{P} = (P_1, \ldots, P_n)$. For $\boldsymbol{a} = (a_1, \ldots, a_d) \in [0, 1]^d$, let $S = \prod_{i=1}^{d} [0, a_i)$ be the anchored box $[0, \boldsymbol{a})$, and let $A(\mathcal{P}, S)$ denote the number of points P_i in S. The signed discrepancy of $\mathcal{P}$ at S is

$$\delta(S) = \delta(S; \mathcal{P}) = \frac{A(\mathcal{P}, S)}{n} - \lambda_d(S)$$

and the star-discrepancy of $\mathcal{P}$ is

$$D_n^*(\mathcal{P}) = \sup_{\boldsymbol{a} \in [0,1)^d} |\delta([0, \boldsymbol{a}); \mathcal{P})|\,. \tag{3.1}$$

The significance of the star discrepancy comes from the Koksma-Hlawka inequality:

$$|\hat{\mu} - \mu| \leq D_n^*(\mathcal{P}) V_{\mathrm{HK}}(f) \tag{3.2}$$

where $V_{\mathrm{HK}}(f)$ is the total variation of f in the sense of Hardy and Krause (Niederreiter, 1992).

We can trivially get a small $D_n^*(\mathcal{P})$ by taking x_i to be the preimage under H of a low discrepancy point set in $[0,1]^d$. The Hilbert curve clearly adds no value for such a construction. For practical purposes we consider only x_i generated as low discrepancy points in $[0,1]$.

One such construction is the lattice,

$$x_i = \frac{i-1}{n}. \tag{3.3}$$

The lattice (3.3) has star discrepancy $1/n$ and the lowest possible star discrepancy (Niederreiter, 1992) for n points in $[0,1]$ is $1/(2n)$ attained via $x_i = (i-1/2)/n$ for $i = 1, \ldots, n$.

Another such construction is the van der Corput sequence (van der Corput, 1935). In van der Corput sampling of $[0,1)$, the integer $i \geq 0$ is written in integer base $b \geq 2$ as $i = \sum_{k=1}^{\infty} d_k b^{k-1}$ for $d_k = d_k(i) \in \{0, 1, \ldots, b-1\}$. Then i is mapped to

$$x_i = \sum_{k=1}^{\infty} d_k b^{-k}. \tag{3.4}$$

The star-discrepancy of the van der Corput sequence is $O(n^{-1}\log(n))$. The van der Corput sequence can be extended one point at a time, while the lattice sequence is not extensible except by doubling the sample size. Figures 2 and 3 show the Hilbert mappings from lattice sequence and van der Corput sequence, respectively. For $n = b^m$, the van der Corput sequence in base b is a permutation of the lattice sequence.

In Theorem 3.1 we bound the star-discrepancy of stratified x_i like (3.3). Figure 4 shows some of these strata for small n.

Theorem 3.1. *Let $x_1, \ldots, x_n \in [0,1]$ and let $\mathcal{P} = (P_1, \ldots, P_n)$ where $P_i = H(x_i) \in [0,1]^d$. If each interval $I_k = [(k-1)/n, k/n)$, for $k = 1, \ldots, n$ contains precisely one of the x_i, then*

$$D_n^*(\mathcal{P}) \leq 4d\sqrt{d+3}\, n^{-1/d} + O(n^{-2/d}). \tag{3.5}$$

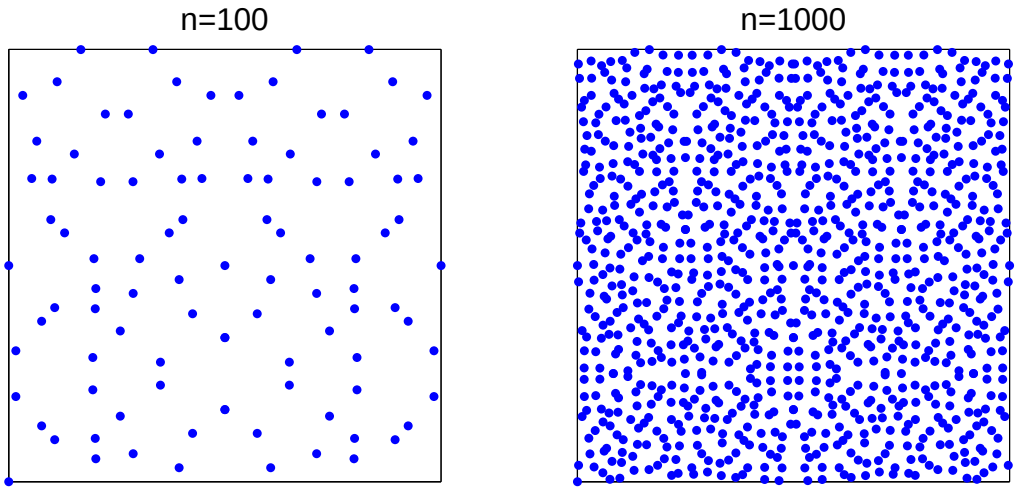

Figure 2: Hilbert mappings of $(i-1)/n$ to $[0,1]^2$ for $i = 1, \ldots, n$ where $n \in \{100, 1000\}$.

Proof. Choose any $\boldsymbol{a} \in [0,1]^d$ and let $S = [0, \boldsymbol{a})$. Next, define $E_k = H(I_k)$ for $k = 1, \ldots, n$ and adjoin $E_{n+1} = H(\{1\})$. By additivity of signed discrepancy,

$$\delta(S; \mathcal{P}) = \frac{1}{n} \sum_{k=1}^{n+1} \delta(S \cap E_k; \mathcal{P}) = \frac{1}{n} \sum_{k=1}^{n} \delta(S \cap E_k; \mathcal{P}),$$

because $S \cap E_{n+1}$ has volume 0 and has no points of $\mathcal{P}$. From here on, we restrict attention to E_k for $k = 1, \ldots, n$. If $E_k \cap S = \emptyset$, then $\delta(S \cap E_k; \mathcal{P}) = 0$. By the measure preserving property of H, $\lambda_d(E_k) = 1/n$, and so if $E_k \subseteq S$, then $\delta(S \cap E_k; \mathcal{P}) = 0$. Otherwise $-1/n \leq \delta(S \cap E_k; \mathcal{P}) \leq 1/n$. Let B be the number of 'boundary' E_k, which intersect both S and $S^c = [0,1]^d \setminus S$. Then $|\delta(S; \mathcal{P})| \leq B/n$, and we turn to bounding B.

Let r_k be the diameter of E_k. By (2.3), $r_k \leq \varepsilon \equiv 2\sqrt{d+3}n^{-1/d}$. Define $S_+ = \prod_{j=1}^d [0, \min(a_j + \varepsilon, 1))$ and $S_- = \prod_{j=1}^d [0, \max(a_j - \varepsilon, 0))$. If E_k intersects S and S^c, then $E_k \subset S_+ \setminus S_-$. Because the E_k are disjoint with volume $1/n$,

$$B \leq n\lambda_d(S_+ \setminus S_-) \leq 2n\left(\prod_{j=1}^d (a_j + \varepsilon) - \prod_{j=1}^d a_j \right) \leq 2n(d\varepsilon + O(\varepsilon^2)).$$

Thus $|\delta(S; \mathcal{P})| \leq 2d\varepsilon + O(\varepsilon^2) = 4d\sqrt{d+3}n^{-1/d} + O(n^{-2/d})$, and since S was

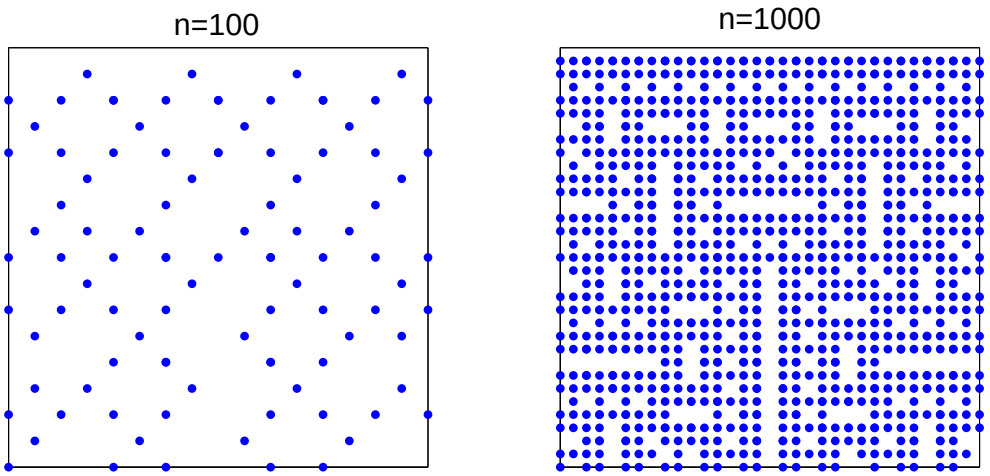

Figure 3: Hilbert mappings of the first n van der Corput points in base 2 to $[0,1]^2$, for $n \in \{100, 1000\}$.

Figure 4: Uniform partitions of $[0,1]^2$ by the mapping H for $n = 3, 8$.

any anchored box, the result now follows. □

In Theorem 3.2 we apply Theorem 3.1 to get a bound for star discrepancy of the van der Corput sequence when $d > 1$. The case $d = 1$ is well known and has star discrepancy $O(\log(n)/n)$.

Theorem 3.2. *For integer base $b \geq 2$ and $n \geq 1$, let $x_1, \ldots, x_n \in [0, 1)$ be defined by the van der Corput mapping (3.4) and let $\mathcal{P} = (P_1, \ldots, P_n)$ where $P_i = H(x_i) \in [0, 1]^d$. Then, for $d > 1$,*

$$D_n^*(\mathcal{P}) \leq \frac{4(b-1)\sqrt{d+3}}{1 - b^{-(d-1)/d}}\, n^{-1/d} + O(n^{-2/d}\log(n)). \tag{3.6}$$

Proof. We begin by writing $n = \sum_{j=0}^{k} a_j b^j$ where $a_j \in \{0, 1, \ldots, b-1\}$ and $a_k > 0$. The x_i can be partitioned into disjoint sets $\mathcal{X}_{j\ell}$ of length b^j for $\ell = 1, \ldots, a_j$. Each of these sets satisfies the conditions of Theorem 3.1. Let $\mathcal{P}_{j\ell}$ be the image of the points in $\mathcal{X}_{j\ell}$ under H.

Now let S be any anchored box $[0, \boldsymbol{a}) \subset [0, 1]^d$. By additivity of local discrepancy over samples, $n\delta(S; \mathcal{P}) = \sum_{j=0}^{k} \sum_{\ell=1}^{a_j} b^j \delta(S; \mathcal{P}_{j\ell})$. Therefore

$$\begin{aligned} n|\delta(S;\mathcal{P})| &\leq \sum_{j=0}^{k} \sum_{\ell=1}^{a_j} b^j \left[4d\sqrt{d+3}b^{-j/d} + O(b^{-2j/d})\right] \\ &\leq C \sum_{j=0}^{k} b^j \left[b^{-j/d} + O(b^{-2j/d})\right] \end{aligned}$$

for $C = (b-1)4d\sqrt{d+3}$. Now for $d > 1$,

$$\sum_{j=0}^{k} (b^{1-1/d})^j \leq \sum_{j=-\infty}^{k} (b^{1-1/d})^j = \frac{(b^{1-1/d})^k}{1 - b^{-1+1/d}} \leq \frac{n^{1-1/d}}{1 - b^{-1+1/d}}$$

and $\sum_{j=0}^{k} (b^{(1-2/d)})^j \leq k b^{k(1-2/d)} = O(n^{1-2/d}\log(n))$. □

Theorems 3.1 and 3.2 show that the star-discrepancy is $O(n^{-1/d})$ for the two sequences. Thus the estimate (2.4) has a worse upper bound than ordinary QMC if the integrand is of bounded variation in the sense of Hardy and Krause.

4 Randomization

In this section, we study the variance resulting from randomized samples along the Hilbert curve. We get convergence rates for Lipschitz continuous functions.

We also study discontinuous functions of the form $f(x) = g(x)1_\Omega(x)$ where the set $\Omega \subset [0,1]^d$ has a boundary that admits $(d-1)$-dimensional Minkowski content (defined below). Functions of this type typically have infinite variation in the sense of Hardy and Krause (Owen, 2005) unless the set Ω is an axis parallel box (or finite union of such). Infinite variation renders the Koksma-Hlawka inequality (3.2) useless. We do know that if $f \in L^2[0,1]^d$ then $f \circ H \in L^2[0,1]$ and scrambled net quadrature on $[0,1]$ for $f \circ H$ will have a mean squared error $o(n^{-1})$. Here we find a rate.

4.1 Randomized Lattice Sequence

We randomize the lattice points in (3.3) by performing a random shift in each subinterval, that is

$$x_i = \frac{i - 1 + \Delta_i}{n}, \quad \text{with} \quad \Delta_i \overset{\text{iid}}{\sim} \mathbb{U}([0,1]), \quad i = 1, \dots, n. \tag{4.1}$$

As a result, $x_i \sim \mathbb{U}(I_i)$ independently for $I_i = [\frac{i-1}{n}, \frac{i}{n}]$. Let $\Delta = (\Delta_1, \dots, \Delta_n)$. A randomized version of (2.4) is given by

$$\hat{\mu}(\Delta) = \frac{1}{n} \sum_{i=1}^{n} f(H(x_i)). \tag{4.2}$$

First, we need some definitions.

Definition 4.1. For a function $f(x)$ defined on $[0,1]^d$, if there exists a constant M such that

$$|f(x) - f(y)| \leq M \|x - y\|$$

for any $x, y \in [0,1]^d$, then $f(x)$ is said to be Lipschitz continuous.

Definition 4.2. For a set $\Omega \subset [0,1]^d$, define

$$\mathcal{M}(\partial\Omega) = \lim_{\epsilon \downarrow 0} \frac{\lambda_d((\partial\Omega)_\epsilon)}{2\epsilon},$$

where $(\partial\Omega)_\epsilon = \left\{ X \in \mathbb{R}^d \mid \text{dist}(x, \partial\Omega) \leq \epsilon \right\}$. If $\mathcal{M}(\partial\Omega)$ exists and is finite, then $\partial\Omega$ is said to admit $(d-1)$-dimensional Minkowski content.

Theorem 4.3. *The estimate $\hat{\mu}(\Delta)$ from (4.2) is unbiased for any $f \in L^2([0,1]^d)$. If f is Lipschitz continuous, then*

$$Var(\hat{\mu}(\Delta)) = O(n^{-1-2/d}). \tag{4.3}$$

If $f(x) = g(x)1_\Omega(x)$ where h is Lipschitz continuous and $\partial\Omega$ admits $(d-1)$-dimensional Minkowski content, then

$$Var(\hat{\mu}(\Delta)) = O(n^{-1-1/d}). \tag{4.4}$$

Proof. Let $E_i = H(I_i)$. Because $x_i \sim \mathbb{U}(I_i)$, we have $H(x_i) \sim \mathbb{U}(E_i)$. Moreover $\lambda_d(E_i) = 1/n$. Thus $\mathbb{E}[\hat{\mu}(\Delta)]$ equals

$$\frac{1}{n}\sum_{i=1}^n \mathbb{E}[f(H(x_i))] = \frac{1}{n}\sum_{i=1}^n \left(n \int_{H(I_i)} f(x)\,\mathrm{d}x \right) = \int_{[0,1]^d} f(x)\,\mathrm{d}x = \mu,$$

and so (4.2) is unbiased.

Let f be a Lipschitz continuous function, and let M be the constant from Definition 4.1. For any $x, y \in E_i$, we have $|f(x) - f(y)| \le M r_i$, where r_i is the diameter of E_i. As in the proof of Theorem 3.1, $r_i \le \varepsilon \equiv 2\sqrt{d+3}n^{-1/d}$, and so $|f(x) - f(y)| \le 2M\sqrt{d+3}n^{-1/d}$. It follows that

$$|f(H(x_i)) - \mathbb{E}[f(H(x_i))]| \le M\varepsilon = 2M\sqrt{d+3}n^{-1/d}, \quad i = 1, \dots, n.$$

Now, since $H(x_i)$'s are independent,

$$\mathrm{Var}(\hat{\mu}(\Delta)) \le \frac{1}{n^2}\sum_{i=1}^n 4M^2(d+3)n^{-2/d} = 4M^2(d+3)n^{-1-2/d}$$

establishing (4.3).

Next consider $f(x) = g(x)1_\Omega(x)$. Let $\mathcal{T}_{\text{int}} = \{1 \le i \le n \mid E_i \subset \Omega\}$ and $\mathcal{T}_{\text{bdy}} = \{1 \le i \le n \mid E_i \cap \Omega \ne \emptyset\} \setminus \mathcal{T}_{\text{int}}$. These are, respectively, the collections of E_i that are interior to Ω, and at the boundary of Ω. Then

$$\hat{\mu}(\Delta) = \frac{1}{n}\sum_{i \in \mathcal{T}_{\text{int}}} g(H(x_i)) + \frac{1}{n}\sum_{i \in \mathcal{T}_{\text{bdy}}} g(H(x_i))1_\Omega(H(x_i)) = \hat{\mu}_{\text{int}} + \hat{\mu}_{\text{bdy}}.$$

Since $g(x)$ is Lipschitz continuous, $\mathrm{Var}(\hat{\mu}_{\text{int}}) = O(n^{-1-2/d})$ by the reasoning above. Also, there exists a constant D with $|g(x)| \le D$ for all $x \in [0,1]^d$. Thus

$$\mathrm{Var}(\hat{\mu}_{\text{bdy}}) = \frac{1}{n^2}\sum_{i \in \mathcal{T}_{\text{bdy}}} \mathrm{Var}(g(H(x_i))1_\Omega(H(x_i))) \le \frac{D^2\,|\mathcal{T}_{\text{bdy}}|}{n^2}. \tag{4.5}$$

Recall that $\partial\Omega$ admits $(d-1)$-dimensional Minkowski content. It follows from Definition 4.2 that

$$\mathcal{M}(\partial\Omega) = \lim_{\epsilon \downarrow 0} \frac{\lambda_d((\partial\Omega)_\epsilon)}{2\epsilon} < \infty.$$

Thus for any fixed $\delta > 2$, there exists $\epsilon_0 > 0$ such that $\lambda_d((\partial\Omega)_\epsilon) < \delta\mathcal{M}(\partial\Omega)\epsilon$ whenever $\epsilon < \epsilon_0$. We can assume that $n > (2\sqrt{d+3}/\epsilon_0)^d$. Then $r_i \le \varepsilon < \epsilon_0$. Notice that $\bigcup_{i\in\mathcal{T}_{\text{bdy}}} E_i \subset (\partial\Omega)_\varepsilon$. We thus arrive at

$$|\mathcal{T}_{\text{bdy}}| \le \frac{\lambda_d((\partial\Omega)_\varepsilon)}{\lambda_d(E_i)} \le \frac{\delta\mathcal{M}(\partial\Omega)\varepsilon}{n^{-1}} = 2\sqrt{d+3}\delta\mathcal{M}(\partial\Omega)n^{1-1/d}.$$

Now by (4.5), we have $\text{Var}(\hat{\mu}_{\text{bdy}}) = O(n^{-1-1/d})$. Finally, from $\text{Var}(\hat{\mu}(\Delta)) \le (\sqrt{\text{Var}(\hat{\mu}_{\text{int}})} + \sqrt{\text{Var}(\hat{\mu}_{\text{bdy}})})^2$, we obtain $\text{Var}(\hat{\mu}(\Delta)) = O(n^{-1-1/d})$. □

Remark 4.4. If Ω is a convex set, then it is easy to see that $\partial\Omega$ admits $(d-1)$-dimensional Minkowski content. Moreover, $\mathcal{M}(\partial\Omega) \le 2d$ as the outer surface area of a convex set in $[0,1]^d$ is bounded by the surface area of the unit cube $[0,1]^d$, which is $2d$. Generally, Ambrosio et al. (2008) show that if Ω has Lipschitz boundary, then $\partial\Omega$ admits $(d-1)$-dimensional Minkowski content. In their terminology, a set Ω is said to have Lipschitz boundary if for every boundary point a there exists a neighborhood A of a, a rotation R in $\mathbb{R}^d$ and a Lipschitz function $f : \mathbb{R}^{d-1} \to \mathbb{R}$ such that $R(\Omega \cap A) = \{(x,y) \in (\mathbb{R}^{d-1} \times \mathbb{R}) \cap R(A) | y \ge f(x)\}$. In other words, $\Omega \cap A$ is the epigraph of a Lipschitz function.

Remark 4.5. The convergence rate (4.4) for $f(x) = g(x)1_\Omega(x)$ extends to functions $f(x) = g_0(x) + \sum_{j=1}^J g_j(x)1_{\Omega_j}(x)$ where all of the g_j are Lipschitz continuous and all of the Ω_j have boundaries with finite Minkowski content.

4.2 Randomized van der Corput Sequence

For the van der Corput sequence, we apply the nested uniform digit scrambling of Owen (1995). Let $a_1, \ldots, a_n$ be the first n points of van der Corput sequence in base b. We may write a_i in base b expansion $a_i = \sum_{j=1}^\infty a_{ij}b^{-j}$, where $0 \le a_{ij} < b$ for all i, j. The scrambled version of $a_1, \ldots, a_n$ is a sequence $x_1, \ldots, x_n$ written as $x_i = \sum_{j=1}^\infty x_{ij}b^{-j}$, where x_{ij} are defined in terms of random permutations of the a_{ij}. The permutation applied to a_{ij} depends

on the values of a_{ih} for $h < j$. Specifically $x_{i1} = \pi(a_{i1})$, $x_{i2} = \pi_{a_{i1}}(a_{i2})$, $x_{i3} = \pi_{a_{i1}a_{i2}}(a_{i3})$, and generally

$$x_{ij} = \pi_{a_{i1}a_{i2}\cdots a_{ij-1}}(a_{ij}).$$

Each permutation $\pi_\bullet$ is uniformly distributed over the $b!$ permutations of $\{0, 1, \ldots, b-1\}$, and the permutations are mutually independent. Let Π be the collection of all the permutations involved in the scrambling scheme. The randomized version of (2.4) becomes

$$\hat{\mu}(\Pi) = \frac{1}{n}\sum_{i=1}^{n} f(H(x_i)). \tag{4.6}$$

Owen (1995) shows that each x_i is uniformly distributed on $[0, 1]$. Thus the estimate (4.6) is unbiased. Moreover, if $n = b^m$ for some nonnegative m, then we can reorder the data values in scrambled sequence such that $x_i \sim \mathbb{U}([\frac{i-1}{n}, \frac{i}{n}])$ independently for $i = 1, \ldots, b^m$. In this case, the scrambled van der Corput sequence is the same as the randomized lattice sequence. Thus the estimate (4.6) has the same variance shown in Theorem 4.3. For an arbitrary sample size n, we can find the associated rates by exploiting the properties of van der Corput sequences.

Theorem 4.6. *The estimate $\hat{\mu}(\Pi)$ of (4.6) is unbiased for any $f \in L^2([0, 1]^d)$. If f is Lipschitz continuous, then*

$$Var(\hat{\mu}(\Pi)) = \begin{cases} O(n^{-1-2/d}), & d \geq 3 \\ O(n^{-2}\log(n)^2), & d = 2 \\ O(n^{-2}), & d = 1. \end{cases} \tag{4.7}$$

If $f(x) = g(x)1_\Omega(x)$ where $g(x)$ is Lipschitz continuous and $\partial\Omega$ admits $(d-1)$-dimensional Minkowski content, then

$$Var(\hat{\mu}(\Pi)) = \begin{cases} O(n^{-1-1/d}), & d \geq 2 \\ O(n^{-2}\log(n)^2), & d = 1. \end{cases} \tag{4.8}$$

Proof. As in the proof of Theorem 3.2 we may write $n = \sum_{j=0}^{k} a_j b^j$ with $a_j \in \{0, 1, \ldots, b-1\}$ where $a_k > 0$, and split the points into $\sum_{j=0}^{k} a_j$ non-overlapping randomized van der Corput sequences, of which a_j have sample

size b^j. Theorem 4.3 gives variance bounds of the form $Cn^{-1-\alpha}$ where $\alpha = 2/d$ when f is Lipschitz continuous and $\alpha = 1/d$ when $f(x) = g(x)1_\Omega(x)$ for Lipschitz continuous g and $\partial\Omega$ with bounded Minkowski content.

In either case, write $n\hat{\mu}(\Pi) = \sum_{j=0}^{k}\sum_{\ell=1}^{a_j} b^j\hat{\mu}_{j\ell}$, for $\hat{\mu}_{j\ell} = \hat{\mu}_{j\ell}(\Pi)$. Then an elementary inequality based on $|\mathrm{Corr}(\hat{\mu}_{j\ell}, \hat{\mu}_{j'\ell'})| \le 1$ yields

$$\mathrm{Var}(\hat{\mu}(\Pi))^{1/2} \le C^{1/2}\sum_{j=0}^{k}\sum_{\ell=0}^{a_j}\frac{b^j}{n}\mathrm{Var}(\hat{\mu}_{j\ell})^{1/2} \le \frac{(b-1)C^{1/2}}{n}\sum_{j=0}^{k} b^{j(1-\alpha)/2}.$$

Now for $\alpha < 1$

$$\sum_{j=0}^{k} b^{j(1-\alpha)/2} \le \frac{b^{k(1-\alpha)/2}}{1-b^{(\alpha-1)/2}} \le \frac{n^{(1-\alpha)/2}}{1-b^{(\alpha-1)/2}}$$

using $b^k \le n$. Then $\mathrm{Var}(\hat{\mu}(\Pi)) = O(n^{-1-\alpha})$. That is, for $\alpha < 1$, the van der Corput construction inherits the rate of the stratified one.

For $\alpha = 1$ we have $\sum_{j=0}^{k} b^{j(1-\alpha)/2} = k+1 = O(\log(n))$ and then $\mathrm{Var}(\hat{\mu}(\Pi)) = O(n^{-2}\log(n)^2)$. For $\alpha > 1$ we have $\sum_{j=0}^{k} b^{j(1-\alpha)/2} = O(1)$ and then $\mathrm{Var}(\hat{\mu}(\Pi)) = O(n^{-2})$.

Equation (4.7) now follows because $\alpha < 1$ for $d \ge 3$, $\alpha = 1$ for $d = 2$ and $\alpha > 1$ for $d = 1$. Similarly, (4.8) follows because $\alpha < 1$ for $d \ge 2$ and $\alpha = 1$ for $d = 1$. □

Scrambled net quadrature has a mean squared error of $O(n^{-3}\log(n)^{d-1})$ for integrands whose mixed partial derivative taken once with respect to all components of x is in $L^2[0,1]^d$ (Owen, 1997, 2008). The rate in (4.7) for $d = 1$ is not as good as that rate even though the algorithms match in this case. The explanation is that Lipschitz continuity is a weaker condition than having the mixed partial in L^2.

4.3 Adaptive sampling

Integration of discontinuous functions is an important challenge because there are few good solutions for them. From the proof of Theorem 4.3, we see that intervals of $[0,1]$ in which $f \circ H$ is discontinuous contribute $O(n^{-1-1/d})$ to the variance, while the other intervals contribute only $O(n^{-1-2/d})$. This suggests that we might improve matters by oversampling the intervals of discontinuity. In that proof $\mathcal{T}_{\mathrm{bdy}}$ collects the indices of E_i touching the boundary

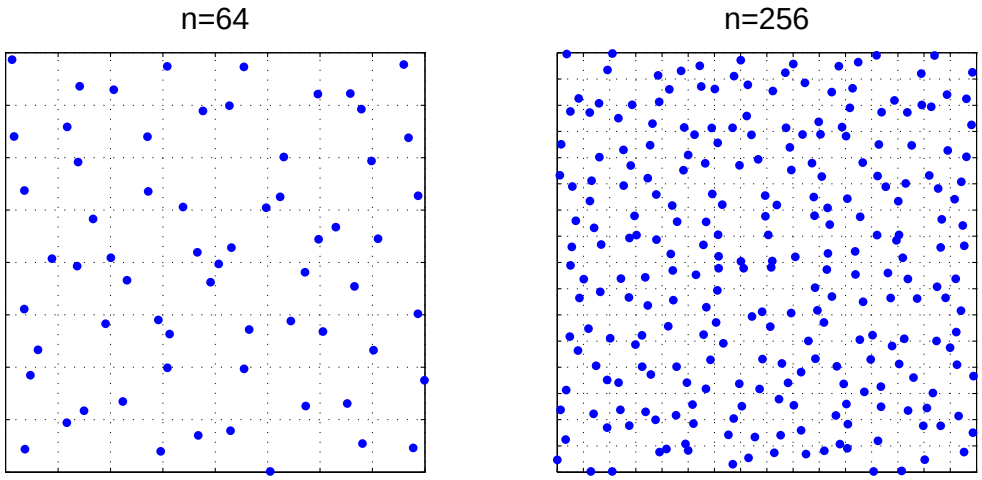

Figure 5: Hilbert mappings into $[0,1]^2$ of n scrambled van der Corput points in base 2, for $n \in \{64, 256\}$.

of the discontinuity, $\mathcal{T}_{\text{int}}$ collects those with $E_i \subset \Omega$ and the ones contained in Ω^c don't contribute to the error. Let us write the estimated integral of the discontinuous function $f(x) = g(x)1_\Omega(x)$ as

$$\hat{\mu} = \frac{1}{n}\sum_{i\in\mathcal{T}_{\text{bdy}}} f(H(x_i)) + \frac{1}{n}\sum_{i\notin\mathcal{T}_{\text{bdy}}} f(H(x_i)). \tag{4.9}$$

Suppose we have prior knowledge about the set $\mathcal{T}_{\text{bdy}}$. We could then use that knowledge to sample $n_0 = \lceil n/|\mathcal{T}_{\text{bdy}}|\rceil$ times in each stratum E_i for $i \in \mathcal{T}_{\text{bdy}}$, and use one sample in the remaining strata as usual. From such samples we get the unbiased estimator

$$\hat{\mu} = \frac{1}{nn_0}\sum_{j=1}^{n_0}\sum_{i\in\mathcal{T}_{\text{bdy}}} f(H(x_i^{(j)})) + \frac{1}{n}\sum_{i\notin\mathcal{T}_{\text{bdy}}} f(H(x_i)), \tag{4.10}$$

where $x_i^{(j)}, x_i \sim \mathbb{U}([\frac{i-1}{n}, \frac{i}{n}])$ independently. The cost of the estimate (4.10) is at most two times the original estimate (4.9) as it makes at most $2n$ function evaluations. Roughly half of the evaluations are in the boundary strata.

Theorem 4.7. *Suppose $f(x) = g(x)1_\Omega(x)$ satisfying the conditions of the second part of Theorem 4.3. Then the variance of $\hat{\mu}$ in (4.10) is $O(n^{-1-2/d})$.*

Proof. From the proof of Theorem 4.3, we have $|\mathcal{T}_{\mathrm{bdy}}| = O(n^{1-1/d})$ and

$$\mathrm{Var}\left(\frac{1}{n}\sum_{i\notin\mathcal{T}_{\mathrm{bdy}}} f(H(x_i))\right) = O(n^{-1-2/d}).$$

It remains to bound the variance of first term in the right side of (4.10). Similarly to (4.5), we find that

$$\begin{aligned}\mathrm{Var}\left(\frac{1}{nn_0}\sum_{j=1}^{n_0}\sum_{i\in\mathcal{T}_{\mathrm{bdy}}} f(H(x_i^{(j)}))\right) &= \frac{1}{n^2n_0^2}\sum_{j=1}^{n_0}\sum_{i\in\mathcal{T}_{\mathrm{bdy}}}\mathrm{Var}(f(H(x_i^{(j)})))\\ &= O(n^{-2}n_0^{-1}\,|\mathcal{T}_{\mathrm{bdy}}|) = O(n^{-1-2/d}),\end{aligned}$$

which completes this proof. □

In practice, we have no prior knowledge of $\mathcal{T}_{\mathrm{bdy}}$ and so Theorem 4.7 describes an unusable method. It does however suggest the possibility of adaptive algorithms that both discover and exploit the presence of boundary intervals.

5 Numerical Study

5.1 Computational Issue

In this section we use the image under H of scrambled van der Corput sampling points on some test integrands with known integrals and compare our observed mean squared errors to the theoretical rates. We chose the van der Corput points in base 2 because it is extensible and is easily expressed in base 2 which conveniently matches the base used to define the Hilbert curve. The first step is to randomize the van der Corput sequence using the scrambling scheme of Owen (1995). In the next step, we use the algorithm given by Butz (1971) for mapping the one-dimensional sequence to a d-dimensional sequence. Butz' algorithm is iterative, requiring a number of iterations equal to the order of the curve, say, m. The accuracy of the approximation of each coordinate is 2^{-m}.

Using Butz' iteration turns an algorithm with n function values into one that costs $O(n\log(n))$. For practical computation, 2^{-m} is set to the machine precision, e.g., $m = 31$ in our numerical examples, thus the effect is negligible.

Suppose we are going to map a point x in $[0,1]$ to d-dimensional point P in $[0,1]^d$, and suppose x is expressed as an md-bit binary number:

$$x = 0._2\rho_1^1\rho_2^1\cdots\rho_d^1\rho_1^2\rho_2^2\cdots\rho_d^2\rho_1^m\rho_2^m\cdots\rho_d^m.$$

Define $\rho^i = 0._2\rho_1^i\rho_2^i\cdots\rho_d^i$. In Butz' algorithm, ρ^i is transformed to $\alpha^i = 0._2\alpha_1^i\alpha_2^i\cdots\alpha_d^i$ via some logical operations. See Butz (1971) for details. The coordinates p_j of P are then given by

$$p_j = 0._2\alpha_j^1\alpha_j^2\cdots\alpha_j^m,$$

for $j = 1,\ldots,d$. To effect this algorithm, one needs to scramble the first md digits of the points in the van der Corput sequence. Suppose that the sample size is $n = 2^k$ for integer $k \geq 0$ with $k < md$. At the scrambling stage, we just need to store $n-1$ permutations to scramble the first k digits. The remaining $md-k$ digits are randomly and independently chosen from $\{0,1\}$. When $md \gg k$, the storage requirement of scrambled van der Corput points is much less than that of scrambling a d-dimensional digital in base 2. Note that the Hilbert computations are very fast since they are based on logical operations.

5.2 Examples

We use three integrands of different smoothness to assess the convergence of our quadrature methods:

- Smooth function: $f_1(X) = \sum_{i=1}^d X_i$;
- Function with cusp: $f_2(X) = \max(\sum_{i=1}^d X_i - \frac{d}{2}, 0)$;
- Discontinuous function: $f_3(X) = 1_{\{\sum_{i=1}^d X_i > \frac{d}{2}\}}(X)$.

Note that f_1 and f_2 are Lipschitz continuous. From Theorem 4.6, the theoretical rate of mean squared error for these two function is $O(n^{-1-2/d})$. The corresponding rate for f_3 is $O(n^{-1-1/d})$ as its discontinuity boundary is has finite Minkowski content. Figure 6 shows the convergence graphs for $d = 2, 3, 8, 16$. These results support the theoretical rates shown in Theorem 4.6.

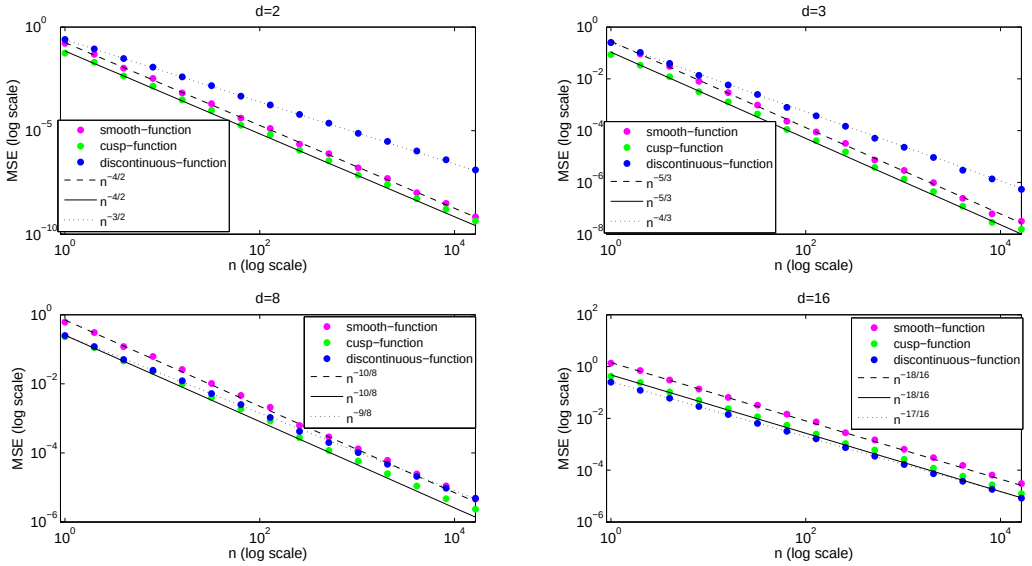

Figure 6: MSE versus n for functions f_1 (smooth) f_2 (cusp) and f_3 (discontinuous), in dimensions $d = 2, 3, 8, 16$. The sample points are first n van der Corput points (scrambled) for $n = 2^k, k = 0, \ldots, 14$. The reference lines are proportional to labeled rates, which reflect the theoretical rates. The MSEs are calculated based on 1000 repetitions.

6 Discussion

In this paper, we study a quadrature method combining the one-dimensional QMC points with the Hilbert curve in dimension d. We find that the star-discrepancy has a very poor convergence rate $O(n^{-1/d})$ in d dimensions, which is the rate one would attain by sampling on an m^d grid. Although this rate seems slow, deterministic quadrature for Lipschitz functions, using $n = m^d$ points in $[0, 1]^d$ has an error rate of $O(n^{-1/d})$ (and a lower bound at that rate) according to Sukharev (1979) as reported in Novak (1988). See also Sobol' (1989). When f has bounded variation on $[0, 1]^d$, then the Hilbert mapping of a low discrepancy point set in $[0, 1]$ attains the optimal rate.

Randomized van der Corput sampling has a mean squared error of $O(n^{-1-2/d})$ for Lipschitz continuous functions. This is the same rate seen for samples of size $n = m^d$ in the stratified sampling method of Dupach (1956) and Haber (1966), which takes one or more points independently in each of m^d congruent subcubes of $[0, 1]^d$. Compared to $O(n^{-1/d})$, this rate reflects the widely

seen error reduction by $O(n^{-1/2})$ commonly seen in the randomized setting versus worst case settings.

Both deterministic and randomized versions of this Hilbert space sampling match the rates seen on grids, without requiring n to be of the form m^d. This is why we think of the Hilbert mapping of van der Corput sequences as extensible grids.

The figures in Gerber and Chopin (2014) show a decreasing rate improvement over Monte Carlo as the dimension of their examples increases. Our results do not yield a convergence rate for their algorithm. They use the inverse $h_m : [0,1]^d \to [0,1]$ of the Hilbert function H_m, in addition to H_m. The function H is not invertible as there is a set of measure 0 in $[0,1]^d$ whose points have more than one pre-image in $[0,1]$. For large m, the function h_m is very non-smooth and has enormous variation because nearby points in $[0,1]^d$ can arise as the images under H_m of widely separated points in $[0,1]$.

Our main theorems 3.1, 3.2, 4.3, and 4.6 on star discrepancy and sampling variance are not strongly tied to the Hilbert space-filling curve. The space-filling curves of Peano and Sierpinski also satisfy the Hölder inequality with exponent $1/d$ that we based our arguments on, although with a different constant. As a result, the same rates of convergence hold for stratified and van der Corput sampling along these curves. The Lebesgue space-filling curve, also called the Z curve, differs from the aforementioned curves in that it is differentiable almost everywhere. It also satisfies Hölder continuity, but the exponent is $\log(2)/(d\log(3))$ which is worse than $1/d$ that holds for the other curves. Using the Lebesgue curve is roughly like multiplying the dimension by $\log_2(3)$, compared to using the Hilbert curve. See Zumbusch (2003, Chapter 4) for these properties of space-filling curves.

Acknowledgments

We thank Erich Novak for helpful comments. ABO was supported by the US NSF under grant DMS-0906056. ZH was supported by a PhD Short-Term Visiting abroad Scholarship of Tsinghua University.

References

Ambrosio, L., Colesanti, A., and Villa, E. (2008). Outer Minkowski content for some classes of closed sets. *Mathematische Annalen*, 342(4):727–748.

Bader, M. (2013). *Space-filling curves: an introduction with applications in scientific computing*, volume 9. Springer, Berlin.

Butz, A. R. (1971). Alternative algorithm for Hilbert's space-filling curve. *IEEE Transactions on Computers*, 20(4):424–426.

Dupach, V. (1956). Stockasticke pocetni metodi. *Casopis pro pestováni matematiky*, 81(1):55–68.

Gerber, M. and Chopin, N. (2014). Sequential quasi-Monte Carlo. *arXiv preprint arXiv:1402.4039*.

Haber, S. (1966). A modified Monte-Carlo quadrature. *Mathematics of Computation*, 20(95):361–368.

L'Ecuyer, P., Lécot, C., and Tuffin, B. (2008). A randomized quasi-Monte Carlo simulation method for Markov chains. *Operations Research*, 56(4):958–975.

McKay, M. D., Beckman, R. J., and Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245.

Netto, E. (1879). Beitrag zur mannigfaltigkeitslehre. *Crelle J*, 86:263–268.

Niederreiter, H. (1992). *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, Philadelphia, PA.

Novak, E. (1988). *Deterministic and stochastic error bounds in numerical analysis*, volume 1349. Springer-Verlag, Berlin.

Owen, A. B. (1995). Randomly permuted (t, m, s)-nets and (t, s)-sequences. In Niederreiter, H. and Shiue, P. J.-S., editors, *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, pages 299–317. Springer.

Owen, A. B. (1997). Scrambled net variance for integrals of smooth functions. *Annals of Statistics*, 25(4):1541–1562.

Owen, A. B. (2005). Multidimensional variation for quasi-Monte Carlo. In Fan, J. and Li, G., editors, *International Conference on Statistics in honour of Professor Kai-Tai Fang's 65th birthday*.

Owen, A. B. (2008). Local antithetic sampling with scrambled nets. *Annals of Statistics*, 36(5):2319–2343.

Rafajłowicz, E. and Skubalska-Rafajłowicz, E. (2008). Equidistributed sequences along space-filling curves in sampling of images. In *16th European signal processing conference*, pages 25–28.

Sagan, H. (1994). *Space-filling curves*, volume 18. Springer-Verlag New York.

Schretter, C. and Niederreiter, H. (2013). A direct inversion method for non-uniform quasi-random point sequences. *Monte Carlo Methods and Applications*, 19(1):1–9.

Sobol', I. (1989). Quadrature formulae for functions of several variables satisfying a general Lipschitz condition. *USSR Computational Mathematics and Mathematical Physics*, 29(3):201–206.

Sukharev, A. G. (1979). Optimal numerical integration formulas for some classes of functions of several variables. *Soviet Math. Doklady*, 20:472–475.

van der Corput, J. G. (1935). Verteilungsfunktionen I. *Nederl. Akad. Wetensch. Proc.*, 38:813–821.

Zumbusch, G. (2003). *Parallel multilevel methods*. B. G. Teubner, Wiesbaden.